

(12) **United States Patent**  
**Zhang et al.**

(10) **Patent No.:** **US 7,620,652 B2**  
(45) **Date of Patent:** **Nov. 17, 2009**

(54) **PROCESSING STRUCTURED DATA**

2002/0038319 A1 3/2002 Yahagi  
2002/0087596 A1 7/2002 Lewontin  
2002/0143521 A1 10/2002 Call

(75) Inventors: **Zhengyu Zhang**, Milpitas, CA (US);  
**Hui Tian**, Milpitas, CA (US)

(73) Assignee: **Ximpleware, Inc.**, Milpitas, CA (US)

#### OTHER PUBLICATIONS

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 464 days.

U.S. Appl. No. 60/378,018, filed May 10, 2002.\*

"Algorithms and programming models for efficient representation of XML for Internet applications", by Sundaresan et al., WWW10, May 1-5, 2001, Hong Kong.\*

"Schema Extraction for Multimedia XML Document Retrieval", by Yoon et al, Proceedings of the First International Conference on Web Information Systems Engineering (WISE'00)-vol. 2, Year of Publication: 2000, ISBN: 0-7695-0577-5-2.\*

"Generic Programming for XML Tools", by Jeurig et al., Institute of Information and Computer Sciences, Utrecht University, The Netherlands, May 27, 2002.\*

"Millau: an encoding format for efficient representation and exchange of XML over the web", by Girardot et al., Computer Networks: The International Journal of Computer and Telecommunications Networking, vol. 33, Issue 1-6 (Jun. 2000), pp. 747-765.\*

(21) Appl. No.: **11/502,908**

(22) Filed: **Aug. 11, 2006**

(65) **Prior Publication Data**

US 2006/0282451 A1 Dec. 14, 2006

#### Related U.S. Application Data

(63) Continuation of application No. 10/272,077, filed on Oct. 15, 2002, now Pat. No. 7,133,857.

(Continued)

(51) **Int. Cl.**

**G06F 7/00** (2006.01)

**G06F 17/00** (2006.01)

*Primary Examiner*—Tim T. Vo

*Assistant Examiner*—Jay A Morrison

(74) *Attorney, Agent, or Firm*—Beyer Law Group LLP

(52) **U.S. Cl.** ..... **707/102; 707/2; 707/4;**  
**707/100; 707/101**

(57) **ABSTRACT**

(58) **Field of Classification Search** ..... **707/102,**  
**707/2, 4, 100, 101**  
See application file for complete search history.


The present invention provides a fast and efficient way of processing structured data by utilizing an intermediate file to store the structural information. The structured data may be processed into a Binary mask Format (BMF) file which may serve as a starting point for post-processing. A tree structure built on top of the BMF file may be constructed very quickly, and also takes up less space than a DOM tree. Additionally, BMF records may reside entirely in the memory and contain structural information, allowing SAX-like sequential data access.

(56) **References Cited**

#### U.S. PATENT DOCUMENTS

5,748,953 A \* 5/1998 Mizutani et al. .... 707/6  
6,209,124 B1 3/2001 Vermeire et al.  
6,418,446 B1 7/2002 Lection et al.  
6,947,932 B2 9/2005 Brandin et al.  
7,127,467 B2 \* 10/2006 Yalamanchi et al. .... 707/102  
7,321,900 B1 1/2008 Karun et al.  
2002/0029229 A1 3/2002 Jakopac et al.

**35 Claims, 6 Drawing Sheets**



**US 7,620,652 B2**

Page 2

---

OTHER PUBLICATIONS

“BitCube: A Three-Dimensional Bitmap Indexing for XML Documents”, by Yoon et al., Journal of Intelligent Information Systems, 17:2/3, 241-254, 2001.\*

“Binary XML Content Format Specification”, Version 1.3, Jul. 25, 2001, Wireless Application Protocol, WAP-192-WBXML-20010725-a.\*

H. Blume et al., Integration of High-Performance ASICs into Reconfigurable Systems Providing Additional Multimedia Functionality, 2002, IEEE, 10 pgs.

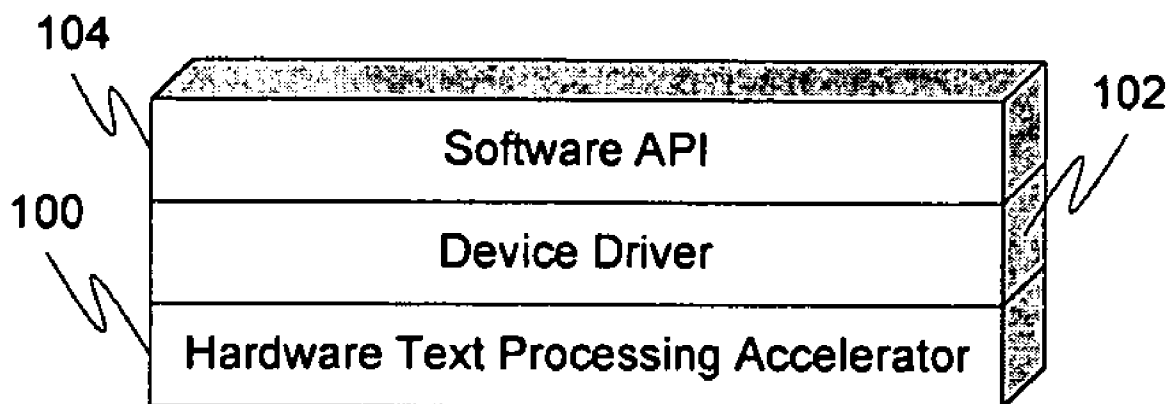
M. Morris Mano, “Computer System Architecture”, Prentice-Hall, Third Edition, 4 pages.

“Efficient Storage of XML Data”, by Kanne et al., In Proceedings of the 16<sup>th</sup> International Conference on Data Engineering, Jun. 16, 1999, 19 Pages.

Office Action dated Sep. 26, 2008 from U.S. Appl. No. 11/581,211.

Office Action dated Mar. 16, 2009 from U.S. Appl. No. 11/581,211.

\* cited by examiner



**FIG. 1**




FIG. 2

U.S. Patent

Nov. 17, 2009

Sheet 3 of 6

US 7,620,652 B2




FIG. 3




FIG. 4

**U.S. Patent**

**Nov. 17, 2009**

**Sheet 5 of 6**

**US 7,620,652 B2**

500	502	504	506	508	510	512	514	516	518	520
<i>N</i>	<i>N</i>	<i>N</i>	<i>N</i>	<i>N</i>	<i>N</i>	<i>N</i>	<i>N</i>	<i>N</i>	<i>N</i>	<i>N</i>
Starting Offset	Depth	Content Type	Modification Indicator	Insertion Indicator	End of Document Indicator	Record in use Indicator	Name space Indicator	Reference Indicator	Reserved	Length

**FIG. 5**

U.S. Patent

Nov. 17, 2009

Sheet 6 of 6

US 7,620,652 B2

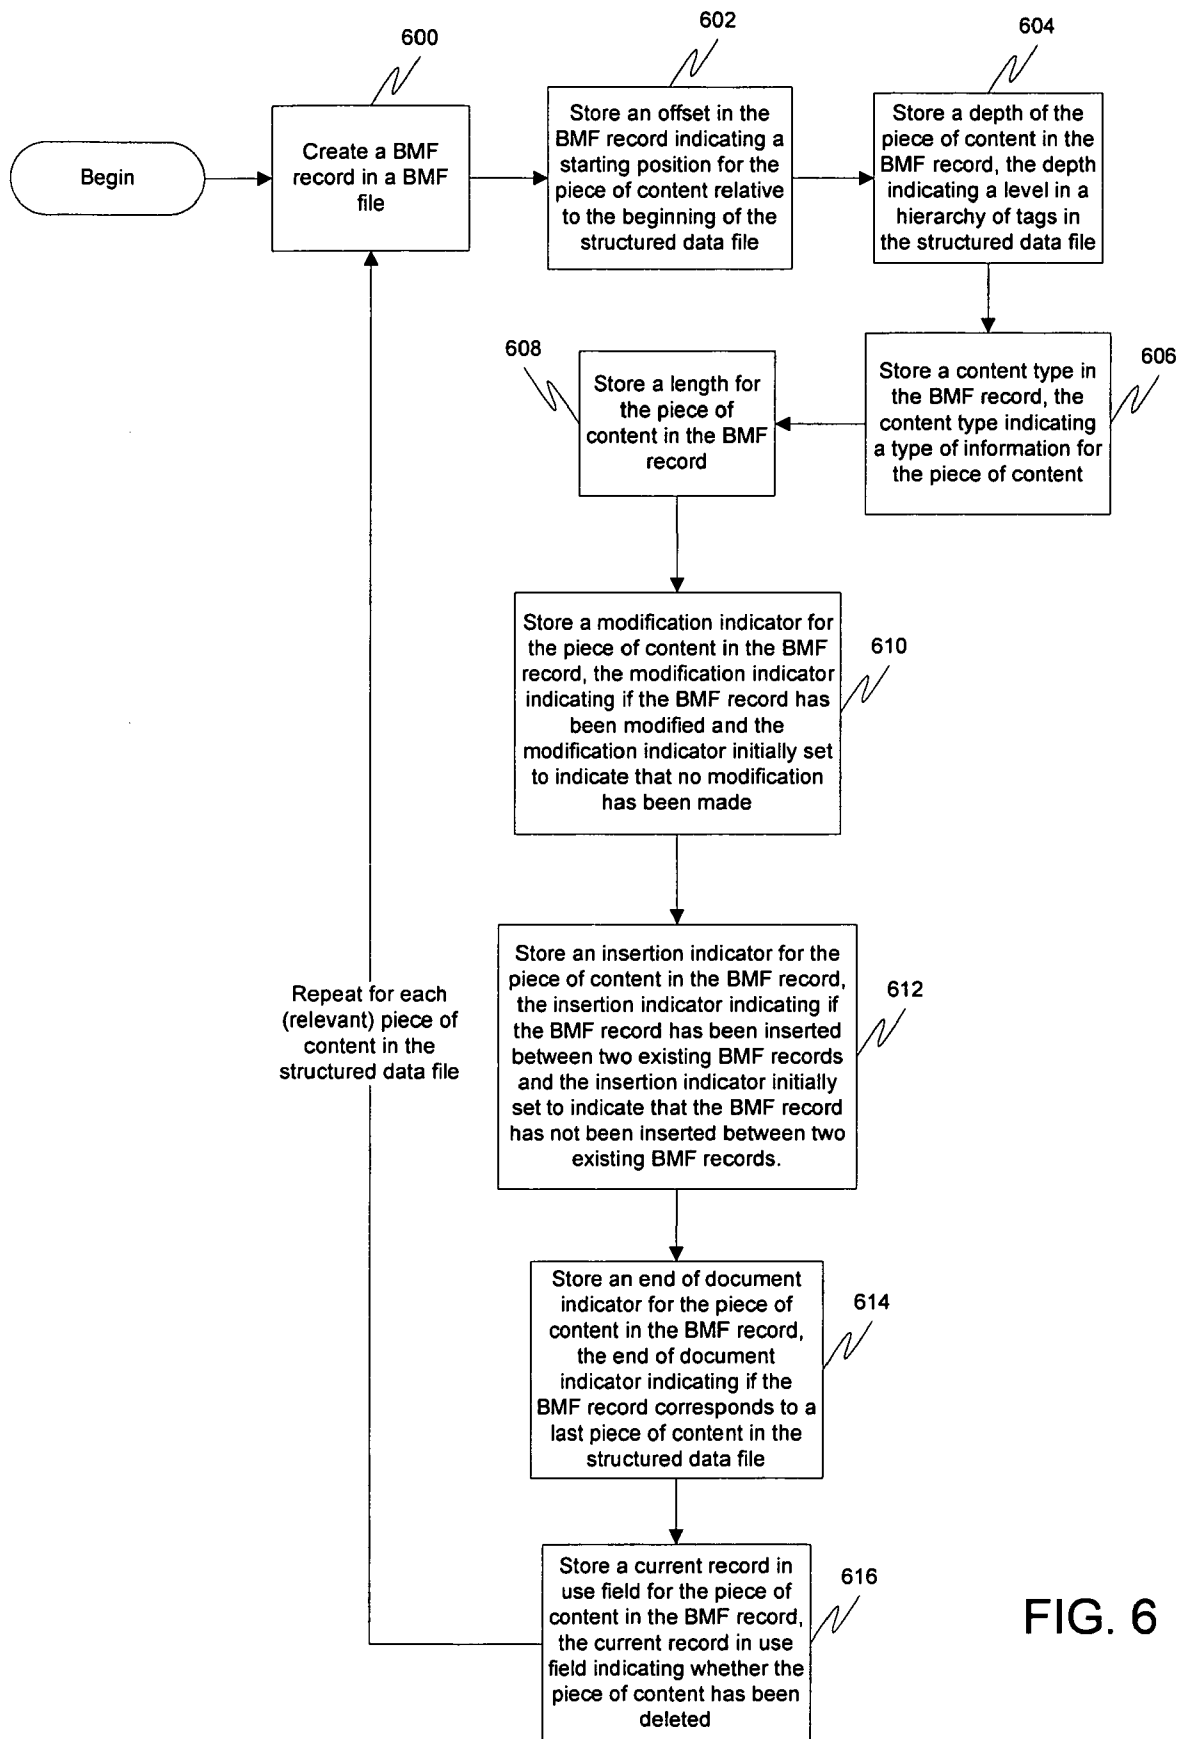


FIG. 6



US 7,620,652 B2

1

**PROCESSING STRUCTURED DATA**

This application is a continuation of U.S. patent application Ser. No. 10/272,077, filed Oct. 15, 2002 now U.S. Pat. No. 7,133,857, entitled "PROCESSING STRUCTURED DATA," which is hereby incorporated by reference.

**FIELD OF THE INVENTION**

The present invention relates to the field of structured data files in computer systems. More specifically, the present invention relates to the processing of structured data in an efficient manner.

**BACKGROUND OF THE INVENTION**

Structured data represents a large portion of the information accessed on the Internet and other computer networks. There are several reasons why structured data is so popular. American Standard Code for Information Interchange (ASCII) and its extensions, such as Unicode Transformation Formats UTF-8 and UTF-16 are among the most common standard encoding formats. Text encoding puts information into a format that is easily readable by a human, thus it is easy for programmers to develop and debug applications. Lastly, textual encoding is extensible and adding new information may be as simple as adding a new key-value pair.

Recently, Extensible Markup Language (XML) has been growing in popularity. XML is a markup language for documents containing structured information. Unlike its predecessor, Hypertext Markup Language (HTML), where tags are used to instruct a web browser how to render data, in XML the tags are designed to describe the data fields themselves. XML, therefore, provides a facility to define tags and the structural relationships between them. This allows a great deal of flexibility in defining markup languages to using information. Because XML is not designed to do anything other than describe what the data is, it serves as the perfect data interchange format.

XML, however, is not without its drawbacks. Compared with other data formats, XML can be very verbose. Processing an XML file can be very CPU and memory intensive, severely degrading overall application performance. Additionally, XML suffers many of the same problems that other software-based text-based processing methods have. Modern processors prefer binary data representations, particularly ones that fit the width of the registers, over text-based representations. Furthermore, the architecture of many general-purpose processors trades performance for programmability, thus making them ill-suited for text processing. Lastly, the efficient parsing of structured text, no matter the format, can present a challenge because of the added steps required to handle the structural elements.

Most current XML parsers are software-based solutions that follow either the Document Object Model (DOM) or Simple API for XML (SAX) technologies. DOM parsers convert an XML document into an in-memory hierarchical representation (known as a DOM tree), which can later be accessed and manipulated by programmers through a standard interface. SAX parsers, on the other hand, treat an XML document as a stream of characters. SAX is event-driven, meaning that the programmer specifies an event that may happen, and if that event occurs, SAX gets control and handles the situation.

In general, DOM and SAX are complementary, not competing, XML processing models, each with its own benefits and drawbacks. DOM programming is programmer-friendly, as the processing phase is separate from application logic. Additionally, because the data resides in the memory, repetitive access is fast and flexible. However, DOM requires that

2

the entire document data structure, usually occupying 7-10 times the size of the original XML document, be loaded into the memory, thus making it impractical for large XML documents. SAX, on the other hand, can be efficient in parsing large XML documents (at least when only small amounts of information need to be processed at once), but it maintains little of the structural information of the XML data, putting more of a burden on programmers and resulting in code that is hardwired, bulky, and difficult to maintain.

What is needed is an application program interface (API) that combines the best attributes of both DOM and SAX parsing.

**BRIEF DESCRIPTION OF THE INVENTION**

The present invention provides a fast and efficient way of processing structured data by utilizing an intermediate file to store the structural information. The structured data may be processed into a Binary mask Format (BMF) file which may serve as a starting point for post-processing. A tree structure built on top of the BMF file may be constructed very quickly, and also takes up less space than a DOM tree. Additionally, BMF records may reside entirely in the memory and contain structural information, allowing SAX-like sequential data access.

**BRIEF DESCRIPTION OF THE DRAWINGS**

The accompanying drawings, which are incorporated into and constitute a part of this specification, illustrate one or more embodiments of the present invention and, together with the detailed description, serve to explain the principles and implementations of the invention.

In the drawings:

FIG. 1 is a block diagram illustrating a layer view of an apparatus for efficiently processing structured data in accordance with an embodiment of the present invention.

FIG. 2 is a block diagram illustrating an apparatus for efficiently processing structured data in accordance with an embodiment of the present invention.

FIG. 3 is a timing diagram illustrating the operation of the hardware in accordance with an embodiment of the present invention.

FIG. 4 is a flow diagram illustrating a method for modifying the content of a target string in a BMF file from an old string to a new string in accordance with an embodiment of the present invention.

FIG. 5 is a diagram illustrating a BMF record format in accordance with an embodiment of the present invention.

FIG. 6 is a flow diagram illustrating a method for efficiently processing a structured data file, the structured data file including one or more pieces of content, in accordance with an embodiment of the present invention.

**DETAILED DESCRIPTION**

Embodiments of the present invention are described herein in the context of a system of computers, servers, and software. Those of ordinary skill in the art will realize that the following detailed description of the present invention is illustrative only and is not intended to be in any way limiting. Other embodiments of the present invention will readily suggest themselves to such skilled persons having the benefit of this disclosure. Reference will now be made in detail to implementations of the present invention as illustrated in the accompanying drawings. The same reference indicators will be used throughout the drawings and the following detailed description to refer to the same or like parts.

In the interest of clarity, not all of the routine features of the implementations described herein are shown and described. It

US 7,620,652 B2

3

will, of course, be appreciated that in the development of any such actual implementation, numerous implementation-specific decisions must be made in order to achieve the developer's specific goals, such as compliance with application- and business-related constraints, and that these specific goals will vary from one implementation to another and from one developer to another. Moreover, it will be appreciated that such a development effort might be complex and time-consuming, but would nevertheless be a routine undertaking of engineering for those of ordinary skill in the art having the benefit of this disclosure.

In accordance with the present invention, the components, process steps, and/or data structures may be implemented using various types of operating systems, computing platforms, computer programs, and/or general purpose machines. In addition, those of ordinary skill in the art will recognize that devices of a less general purpose nature, such as hard-wired devices, field programmable gate arrays (FPGAs), application specific integrated circuits (ASICs), or the like, may also be used without departing from the scope and spirit of the inventive concepts disclosed herein.

For purposes of this disclosure, a structured data file is any file containing content as well as some information regarding the structural organization of the content. The present invention provides a fast and efficient way of processing structured data by utilizing an intermediate file to store the structural information. The structured data may be processed into a Binary mask Format (BMF) file which may serve as a starting point for post-processing. A tree structure built on top of the BMF file may be constructed very quickly, and also takes up less space than a DOM tree. Additionally, BMF records may reside entirely in the memory and contain structural information, allowing SAX-like sequential data access. However, while this document will describe advantages that the present invention provides over DOM or SAX, one of ordinary skill in the art will recognize that the present invention need not be limited to replacing DOM or SAX, and can be expanded to non-XML type processing.

FIG. 1 is a block diagram illustrating a layer view of an apparatus for efficiently processing structured data in accordance with an embodiment of the present invention. The apparatus may comprise three layers. A hardware text processing accelerator **100**, occupying the lowest layer, may offer the horsepower necessary to relieve the central processing unit (CPU) from the most processor intensive part of the task. On top of the hardware text processing accelerator **100** may lie a device driver layer **102** that is responsible for the communication between the hardware text processing accelerator **100** and a software layer **104**. The software layer **104** may be designed to offer maximum flexibility and further improve the performance. It may export APIs that are standard-compliant.

The hardware may be designed such that it may quickly match multiple patterns against an incoming data stream. FIG. 2 is a block diagram illustrating an apparatus for efficiently processing structured data in accordance with an embodiment of the present invention. A text processor **200** may be the core of the accelerator. It may be composed of multiple Finite State Machines (FSMs) that process an incoming document in parallel. The output may be the BMF file. It also may set several result registers (not pictured). A PCI interface **202** may handle all handshaking between the hardware and a server PCI bus **204**. A memory controller **206** may receive commands from the PCI interface **202** and convert the PCI address to on-board memory address space. It also may access the board memory accordingly. Configuration registers **208** may determine the configuration of the text processing pipeline and the organization of the memory controller **206**. It may load default values from configuration ROM **210**. Some of the values may be modified by software through the PCI interface **202**. The Configuration ROM **210**

4

may store the default setting of the text processor configuration. It also may store the configuration map of FPGAs.

A document buffer **212** may store the incoming document. This may be a Synchronous Dynamic Random Access Memory (SDRAM). Paging may be utilized if the incoming document is larger than the total buffer size. A BMF buffer **214** may store the output BMF files, together with several other text processor result register values. This may be a separate SDRAM, although in one embodiment of the present invention it may share a single SDRAM with the document buffer. A string cache **216** may buffer the incoming data to smooth out and speed up SDRAM access. A DMA engine **218** may grab the frame data from server main memory and send it back the BMF file.

FIG. 3 is a timing diagram illustrating the operation of the hardware in accordance with an embodiment of the present invention. A reset **300** may be sent out by the host computer system, or by a specific application. While reset is asserted, the configuration may be loaded from ROM to the configuration registers. The text processor then may be set to a default state according to the configuration registers. When software calls the device driver, a start signal **302** may be sent through the PCI to indicate the beginning of a document processing cycle. Then the PCI master may assert a frame number **304** to indicate the beginning of a write transaction. The PCI master may then drive the address/data **306** to the PCI bus. The PCI target interface may respond, causing the DMA to read the document into the SDRAM document buffer **308**. There may also be certain PCI commands reserved to update the configuration registers. The memory controller may be activated by PCI command. It may start processing data in the string buffer **310**. It also may send sync signals to the memory controller. The transferring and processing may be repeated.

The PCI target may sense a valid window to send data. Then the PCI master may assert the frame number to indicate the beginning of a read transaction **312**. The PCI target holds the bus. The DMA engine may then transfer the BMF and result register data to main memory **314**. When all the data is transferred, the PCI target interface may send an end signal to the device driver **316**. The next document processing cycle may start again with a start signal from the device driver **318**.

The output of the hardware is a BMF. In one embodiment of the present invention, the BMF defines a binary record format that is used to describe various fields in a structured data file. It can be viewed as a two-dimensional field of bits. Each piece of useful information in the structured data file may correspond to a record in the BMF file. A record may comprise a starting offset and length of a target string in the structured data file. It may also comprise the depth value, node type, and bit-wide flags. These will be discussed in more detail below. The total length of a record may be an integer multiple of a 32-bit word—the width of the memory bus in most commercial architectures. Two record types may be defined: a full version of 12 bytes in length, and a compressed version of 8 bytes in length. The full version may be based on the assumption that both the string length and the starting offset are 32-bits wide, whereas the compressed version may assume a 16-bit field, which translates to a maximum length of 65536 bytes.

FIG. 4 is a flow diagram illustrating a method for modifying the content of a target string in a BMF file from an old string to a new string in accordance with an embodiment of the present invention. At **400**, a piece of memory of the length of the new string may be allocated. At **402**, the new string may be filled into the memory. At **404**, a record in the BMF file corresponding to the old string may be located. At **406**, a corresponding bit flag for the record may be changed from relative offset to absolute offset. The first time a file is converted to BMF form, relative offsets may be used. However, as soon as a modification is made to a specific string, it can often be difficult if not impossible to continue to track the

## US 7,620,652 B2

5

relative offset for that string. Therefore, absolute offsets may be utilized for all modified strings. At **408**, an offset value in the BMF record for the old string may be replaced with a pointer value of the new string in memory. At **410**, a length field in the BMF record for the old string may be replaced with the length of the new string.

There are at least three types of possible BMF file modes: read-only mode, read-modify mode, and read-modify-add mode. In read-only mode, records representing various types of nodes may be placed sequentially into a BMF file, leaving no empty records. For example, a leaf-level element may be represented as a record for the starting tag, 2n records (one of r property name and one property value) for n properties, and one record for the text for the element, and finally one record for an ending tag name. The presence of the ending tag record may be used for document validation.

The read-modify mode may be identical to read-only mode except each record allows for limited write-access, meaning content can be altered, but not added.

The read-modify-add mode allows complete write-access, which is done by embedding empty records into the file.

The record format may be picked to efficiently represent the necessary information of the original data structure. It may also be made less efficient on purpose to speed up downstream processing.

FIG. 5 is a diagram illustrating a BMF record format in accordance with an embodiment of the present invention. A control word **500** may be thirty-two bits in total in this embodiment. This may include a depth value **502** of sixteen bits. The depth value may indicate the depth of a tag in the hierarchy of tags in the structured data file. Thus, the first tag in a file will have a depth of zero, whereas if another starting tag appears before an ending tag for the first tag, that second starting tag will have a depth of one. A content type **504** may be provided, which indicates what type of information the content is. In an embodiment of the present invention, the following value/content type pairs may be used in this field:

TABLE 1

Content Types and Corresponding Values		
Content Type Value	Content name	Example
0	Starting Tag	<example>
1	Ending Tag	</example>
2	Property Name	<example property1="this">
3	Property Value	<example property2="that">
4	Text	<example> tastyfruit </example>
5	Comment	<!-- thisisacomment -->
6	Processing Instruction	<? ..... ?>
7	Markup declaration I name	<![CDATA[...<<<>>>...]]>
8	Markup declaration I value	<![CDATA[...<<<>>>...]]>
9	Markup declaration II name	<!ENTITY...>
10	Markup declaration II value	<!ENTITY ...>
11	Entity reference	&example.bib;
12	Property Name Value Pair	<example property1="this">

A modification indicator **506** may also be provided, which indicates whether or not the record has been modified. This is important because, as described above, if the record had been modified, then the offset field will contain the real pointer value, not a relative offset. An indicator **508** may indicate that data was inserted in between two existing records. Once

6

again, this is important in determining how to utilize the offset field. If the insertion indicator is set to 1, it indicates that the offset field contains a pointer to an external piece of memory, one that can be used to add child nodes to the current node. An end of document indicator **510** may indicate whether the tag is the last one in the document. This can be important because in some embodiments, ending tags may be ignored when encoding the BMF file in order to save space. Therefore, the last tag in the BMF file may not correspond to the last tag in the structured data file.

A current record in use field **512** may be used to indicate that a record has been deleted. If the field is set to 0, the record may be safely ignored because it has been deleted. A name space indicator **514** may indicate whether or not there is a name space within the token (which may be represented by an "=" sign). A reference bit **516** may indicate when there is an external reference, such as an "&" in a text string.

There may be one or more reserved bits **518**, which are set aside for future uses. Lastly, a length field **520** may indicate the length of the content.

The BMF file, together with the original data in memory, completely describes the original data and its inherent data structure. Traversing the data structure may be easily accomplished using the BMF records. Higher level applications and processing are therefore facilitated by using the BMF. To make it easily accessible and readily integrated to higher level application and processing, device drivers and an application programming interface (API) may be built on top of the BMF.

FIG. 6 is a flow diagram illustrating a method for efficiently processing a structured data file, the structured data file including one or more pieces of content, in accordance with an embodiment of the present invention. The structured data file may be an extensible markup language file. The process loops through each piece of content. In another embodiment of the present invention, the process loops through each relevant piece of content. Relevancy can be determined by the programmer and may be chosen so as to minimize the amount of space used for a BMF file. At **600**, a BMF record is created in a BMF file, the BMF record corresponding to the piece of content. At **602**, an offset may be stored in the BMF record indicating a starting position for the piece of content relative to the beginning of the structured data file. At **604**, a depth of the piece of content may be stored in the BMF record, the depth indicating a level in a hierarchy of tags in the structured data file. At **606**, a content type of the piece of content may be stored in the BMF record, the content type indicating a type of information for the piece of content. The content type may take many forms, such as a starting tag, ending tag, property name, property value, text, comment, processing instruction, markup declaration name, markup declaration value, external reference, property name pair, etc. At **608**, a length may be stored for the piece of content in the BMF record.

At **610**, a modification indicator for the piece of content may be stored in the BMF record, the modification indicator indicating if the BMF record has been modified and the modification indicator initially set to indicate that no modification has been made. At **612**, an insertion indicator for the piece of content may be stored in the BMF record, the insertion indicator indicating if the BMF record has been inserted between two existing BMF records and the insertion indicator initially set to indicate that the BMF record has not been inserted between two existing BMF records. At **614**, an end of document indicator for the piece of content may be stored in the BMF record, the end of document indicator indicating if the BMF record corresponds to a last piece of content in the structured data file. At **616**, a current record in use field may

## US 7,620,652 B2

7

be stored for the piece of content in the BMF record, the current record in use field indicating whether the piece of content has been deleted.

The following example may be used to illustrate an embodiment of the present invention. One of ordinary skill in the art will recognize that this is merely an example and should not be read to be limiting in any way. Suppose an XML file as follows:

```
<?xml version="1.0" encoding="US-ASCII"?>
<benchmark:database xmlns:benchmark=
"http://example.com/xml/benchmark">
<benchmark:person id="012345">
<benchmark:email name="Name012345" />
<!--Edited with XML spy v4.2 -->
<benchmark:line1>L i n e 1 012345 012345</benchmark:line1>
</benchmark:person>
</benchmark:database>
```

An embodiment of the present invention may ignore ending tags and produce the following BMF file:

Starting offset 32 bit	Depth 5 bit	Token type 4 bit	Modify indicator 1 bit	Insertion Indicator 1 bit	End of Document Indicator 1 bit	Current Record in use 1 bit	Name space indicator 1 bit	Reference indicator	unused	Length 16 bit
2	0	6	0	0	0	1	0	0	0	38
42	0	0	0	0	0	1	1	0	0	18
61	0	2	0	0	0	1	1	0	0	14
78	0	3	0	0	0	1	0	0	0	35
116	1	0	0	0	0	1	1	0	0	16
133	1	2	0	0	0	1	0	0	0	2
137	1	3	0	0	0	1	0	0	0	6
147	2	0	0	0	0	1	1	0	0	15
163	2	2	0	0	0	1	0	0	0	4
169	2	3	0	0	0	1	0	0	0	10
185	1	5	0	0	0	1	0	0	0	25
218	2	0	0	0	0	1	1	0	0	15
234	2	4	0	0	0	1	0	0	0	23
0	0	0	0	0	1	1	0	0	0	0

The packet BMF records are:

00000000000000000000000000000000	00000	0110	0	0	0	1	0	0	0	00100110
00000000000000000000000000000000	00000	0000	0	0	0	1	1	0	0	00010010
00000000000000000000000000000000	00000	0010	0	0	0	1	1	0	0	00001110
00000000000000000000000000000000	00001	0011	0	0	0	1	0	0	0	00100011
00000000000000000000000000000000	00001	0000	0	0	0	1	1	0	0	00010000
00000000000000000000000000000000	00001	0010	0	0	0	1	0	0	0	00000010
00000000000000000000000000000000	00010	0011	0	0	0	1	0	0	0	00000110
00000000000000000000000000000000	00010	0000	0	0	0	1	1	0	0	00001111
00000000000000000000000000000000	00010	0010	0	0	0	1	0	0	0	00000100
00000000000000000000000000000000	00001	0011	0	0	0	1	0	0	0	00001010
00000000000000000000000000000000	00010	0101	0	0	0	1	0	0	0	00011001
00000000000000000000000000000000	00010	0000	0	0	0	1	1	0	0	00001111
00000000000000000000000000000000	00010	0100	0	0	0	1	0	0	0	00010111
00000000000000000000000000000000	00000	0000	0	1	1	1	0	0	0	00000000

While embodiments and applications of this invention have been shown and described, it would be apparent to those skilled in the art having the benefit of this disclosure that many more modifications than mentioned above are possible without departing from the inventive concepts herein. The invention, therefore, is not to be restricted except in the spirit of the appended claims.

8

What is claimed is:

1. A method for efficiently processing a structured data file, the structured data file including one or more pieces of content, comprising:

parsing the structured data file by:

creating a first record in an intermediate file, the first record corresponding to a starting tag in the structured data file and having information regarding the starting tag and a content type field identifying the first record as corresponding to a starting tag;

creating a second record in the intermediate file, the second record corresponding to a piece of content in the structured data file, the second record containing one or more descriptors containing information regarding the content, wherein the second record includes an offset indicating a position for said piece of content relative to a point in said structured data file;

creating a third record in the intermediate file, the third record corresponding to an ending tag in the structured data file and having information regarding the

ending tag and a content type field identifying the third record as corresponding to an ending tag; and formatting said intermediate file in a way that allows data from the structured data file to be accessed using both said intermediate file and the structured data file together.



## US 7,620,652 B2

9

2. The method of claim 1, wherein said creating a second record comprises:

creating a binary mask format (BMF) record in the intermediate file, the BMF record corresponding to said piece of content.

3. The method of claim 1, wherein said intermediate file is a BMF file, wherein the BMF file includes one or more BMF records.

4. The method of claim 2, wherein said creating a second record further comprises:

storing a depth value of said piece of content in said BMF record, said depth value indicating a level in a hierarchy of tags in said structured data file.

5. The method of claim 2, wherein said creating a second record further comprises:

storing a content type of said piece of content in said BMF record, said content type indicating a type of information for said piece of content.

6. The method of claim 2, wherein said creating a second record further comprises:

storing a length for said piece of content in said BMF record.

7. The method of claim 2, wherein said offset indicates a starting position for said piece of content relative to a beginning of the structured data file.

8. The method of claim 4, wherein the formatting includes formatting said record in a way that allows data to be accessed using both said intermediate file and the structured data file without traversing the entire structured data file to determine the depth value.

9. A method for efficiently processing a structured data file, the structured data file including one or more pieces of content, comprising:

parsing the structured data file by:

creating a first record in an intermediate file, the first record corresponding to a starting tag in the structured data file and having information regarding the starting tag and a content type field identifying the first record as corresponding to a starting tag;

creating a second record in the intermediate file, the second record corresponding to a piece of content in the structured data file, the second record containing one or more descriptors containing information regarding the content;

storing an offset in said second record indicating a starting position for said piece of content relative to a beginning of the structured data file;

storing a length for said piece of content in said second record;

creating a third record in the intermediate file, the third record corresponding to an ending tag in the structured data file and having information regarding the ending tag and a content type field identifying the third record as corresponding to an ending tag; and

formatting said record in a way that allows data from the structured data file to be accessed using both said intermediate file and the structured data file together.

10. The method of claim 9, wherein said intermediate file is a BMF file, wherein the BMF file includes one or more BMF records.

11. The method of claim 9, wherein said creating a record further comprises:

storing a depth value of said piece of content in said record, said depth value indicating a level in a hierarchy of tags in said structured data file.

12. The method of claim 9, wherein said creating a record further comprises:

10

storing a content type of said piece of content in said record, said content type indicating a type of information for said piece of content.

13. The method of claim 11, wherein the formatting includes formatting said record in a way that allows data to be accessed using both said intermediate file and the structured data file without traversing the entire structured data file to determine the depth value.

14. An apparatus for efficiently processing structured data, comprising:

a peripheral component interface (PCI) interface;

a direct memory access (DMA) engine coupled to said PCI interface;

a text processor coupled to said PCI interface, the text processor configured to:

parse the structured data file by:

creating a first record in an intermediate file, the first record corresponding to a starting tag in the structured data file and having information regarding the starting tag and a content type field identifying the first record as corresponding to a starting tag;

creating a second record in the intermediate file, the second record corresponding to a piece of content in the structured data file, the second record containing one or more descriptors containing information regarding the content, wherein the second record includes an offset indicating a position for said piece of content relative to a point in said structured data file;

creating a third record in the intermediate file, the third record corresponding to an ending tag in the structured data file and having information regarding the ending tag and a content type field identifying the third record as corresponding to an ending tag; and

formatting said intermediate file in a way that allows data from the structured data file to be accessed using both said intermediate file and the structured data file together;

configuration memory coupled to said text processor and to said PCI interface;

a memory controller coupled to said PCI interface;

BMF memory coupled to said DMA engine, said memory controller, and said text processor;

a document buffer coupled to said DMA engine, said memory controller, and said text processor; and

a string cache coupled to said DMA engine, said memory controller, and said text processor.

15. The apparatus of claim 14, wherein said configuration memory includes:

one or more configuration registers; and

configuration read-only-memory coupled to said one or more configuration registers.

16. The apparatus of claim 14, wherein said PCI interface is configured to handle all handshaking between the apparatus and a server PCI bus.

17. The apparatus of claim 14, wherein said memory controller is configured to receive commands from said PCI interface and convert a PCI address to on-board memory address space.

18. The apparatus of claim 17, wherein said memory controller is further configured to access board memory according to said PCI address.

19. The apparatus of claim 17, wherein said configuration register contains a configuration of a text processing pipeline and organization of said memory controller.

## US 7,620,652 B2

11

20. The apparatus of claim 17, wherein said memory buffer is configured to store an incoming document.

21. The apparatus of claim 17, wherein said memory buffer is Synchronous Dynamic Random Access memory (SDRAM).

22. The apparatus of claim 17, wherein said DMA engine is configured to grab frame data from server main memory and send back a BMF file.

23. An apparatus for efficiently processing structured data, comprising:

a peripheral component interface (PCI) interface;

a direct memory access (DMA) engine coupled to said PCI interface;

a text processor coupled to said PCI interface, the text processor configured to:

create a first record in an intermediate file the first record corresponding to a starting tag in the structured data file and having information regarding the starting tag and a content type field identifying the first record as corresponding to a starting tag;

create a second record in the intermediate file, the second record corresponding to a piece of content in the structured data file, the second record containing one or more descriptors containing information regarding the content;

store a length for said piece of content in said second record;

store an offset in said second record indicating a starting position for said piece of content relative to a beginning of the structured data file;

create a third record in the intermediate file, the third record corresponding to an ending tag in the structured data file and having information regarding the ending tag and a content type field identifying the third record as corresponding to an ending tag; and

format said record in a way that allows data from the structured data file to be accessed using both said intermediate file and the structured data file together; configuration memory coupled to said text processor and to said PCI interface;

a memory controller coupled to said PCI interface;

BMF memory coupled to said DMA engine, said memory controller, and said text processor;

a document buffer coupled to said DMA engine, said memory controller, and said text processor; and

a string cache coupled to said DMA engine, said memory controller, and said text processor.

24. The apparatus of claim 23, wherein said configuration memory includes:

one or more configuration registers; and

configuration read-only-memory coupled to said one or more configuration registers.

25. The apparatus of claim 23, wherein said PCI interface is configured to handle all handshaking between the apparatus and a server PCI bus.

26. The apparatus of claim 23, wherein said memory controller is configured to receive commands from said PCI interface and convert a PCI address to on-board memory address space.

27. The apparatus of claim 26, wherein said memory controller is further configured to access board memory according to said PCI address.

28. The apparatus of claim 26, wherein said configuration register contains a configuration of a text processing pipeline and organization of said memory controller.

29. The apparatus of claim 26, wherein said memory buffer is configured to store an incoming document.

12

30. The apparatus of claim 26, wherein said memory buffer is Synchronous Dynamic Random Access memory (SDRAM).

31. The apparatus of claim 26, wherein said DMA engine is configured to grab frame data from server main memory and send back a BMF file.

32. An apparatus for efficiently processing a structured data file, the structured data file including one or more pieces of content, the apparatus comprising:

means for parsing the structured data file by using:

means for creating a first record in an intermediate file, the first record corresponding to a starting tag in the structured data file and having information regarding the starting tag and a content type field identifying the first record as corresponding to a starting tag;

means for creating a second record in the intermediate file, the second record corresponding to a piece of content in the structured data file, the second record containing one or more descriptors containing information regarding the content wherein the second record includes an offset indicating a position for said piece of content relative to a point in said structured data file;

means for creating a third record in the intermediate file, the third record corresponding to an ending tag in the structured data file and having information regarding the ending tag and a content type field identifying the third record as corresponding to an ending tag; and

means for formatting said intermediate file in a way that allows data from the structured data file to be accessed using both said intermediate file and the structured data file together.

33. An apparatus for efficiently processing a structured data file, the structured data file including one or more pieces of content, the apparatus comprising:

means for parsing the structured data file by using:

means for creating a first record in an intermediate file, the first record corresponding to a starting tag in the structured data file and having information regarding the starting tag and a content type field identifying the first record as corresponding to a starting tag;

means for creating a second record in the intermediate file, the second record corresponding to a piece of content in the structured data file, the second record containing one or more descriptors containing information regarding the content;

means for storing an offset in said second record indicating a starting position for said piece of content relative to a beginning of the structured data file;

means for storing a length for said piece of content in said second record; and

means for creating a third record in the intermediate file, the third record corresponding to an ending tag in the structured data file and having information regarding the ending tag and a content type field identifying the third record as corresponding to an ending tag; and

means for formatting said record in a way that allows data from the structured data file to be accessed using both said intermediate file and the structured data file together.

34. A program storage device readable by a machine, tangibly embodying a program of instructions executable by the machine to perform a method for efficiently processing a structured data file, the structured data file including one or more pieces of content, the method comprising:

## US 7,620,652 B2

13

parsing the structured data file by:

- creating a first record in an intermediate file, the first record corresponding to a starting tag in the structured data file and having information regarding the starting tag and a content type field identifying the first record as corresponding to a starting tag; 5
- creating a second record in the intermediate file, the second record corresponding to a piece of content in the structured data file, the second record containing one or more descriptors containing information regarding the content wherein the second record includes an offset indicating a position for said piece of content relative to a point in said structured data file; 10
- creating a third record in the intermediate file, the third record corresponding to an ending tag in the structured data file and having information regarding the ending tag and a content type field identifying the third record as corresponding to an ending tag; and 15
- formatting said intermediate file in a way that allows data from the structured data file to be accessed using both said intermediate file and the structured data file together. 20

35. A program storage device readable by a machine, tangibly embodying a program of instructions executable by the machine to perform a method for efficiently processing a structured data file, the structured data file including one or more pieces of content, the method comprising: 25

14

parsing the structured data file by:

- creating a first record in an intermediate file, the first record corresponding to a starting tag in the structured data file and having information regarding the starting tag and a content type field identifying the first record as corresponding to a starting tag;
- creating a second record in the intermediate file, the second record corresponding to a piece of content in the structured data file, the second record containing one or more descriptors containing information regarding the content;
- storing an offset in said second record indicating a starting position for said piece of content relative to a beginning of the structured data file;
- storing a length for said piece of content in said second record;
- creating a third record in the intermediate file, the third record corresponding to an ending tag in the structured data file and having information regarding the ending tag and a content type field identifying the third record as corresponding to an ending tag; and
- formatting said record in a way that allows data from the structured data file to be accessed using both said intermediate file and the structured data file together.

\* \* \* \* \*

UNITED STATES PATENT AND TRADEMARK OFFICE

**CERTIFICATE OF CORRECTION**

PATENT NO. : 7,620,652 B2  
APPLICATION NO. : 11/502908  
DATED : November 17, 2009  
INVENTOR(S) : Zhang et al.

Page 1 of 1

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:


On the Title Page:

The first or sole Notice should read --

Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 562 days.

Signed and Sealed this

Twenty-sixth Day of October, 2010

A handwritten signature in black ink that reads "David J. Kappos". The signature is written in a cursive, flowing style.

David J. Kappos

*Director of the United States Patent and Trademark Office*